

EV355226642

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Secure Execution of Program Code

Inventor(s):

Paul England

ATTORNEY'S DOCKET NO. MS1-331USC1

SECURE EXECUTION OF PROGRAM CODE

Technical Field

5 The present invention relates to electronic data processing, and more particularly concerns computer hardware and software for manipulating keys and other secure data so as to prevent their disclosure, even to persons having physical control of the hardware and software.

10

Related Applications

This application is related to co-pending commonly assigned application attorney docket MS1-282US (“System and Method for Authenticating an Operating System and proving a CPU/OS to a Third Party”), application docket 777.206US1 (“Loading and Identifying a Digital Rights Management Operating System”), application docket 777.211US1 (“Key-Based Secure Storage”), and application docket 777.212US1 (“Digital Rights Management”). The disclosures of these applications are hereby incorporated by reference.

Copyright Disclaimer

20 A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever. The following notice applies
25 to the software and data as described below and in the drawing hereto: Copyright © 1998, Microsoft Corporation, All Rights Reserved.

Background

More and more digital content is being delivered online over public networks, such as the Internet. For a client, online delivery improves timeliness, 5 convenience, and allows more sophisticated content. For a publisher, online delivery provides mechanisms for enhanced content and reduces delivery costs. Unfortunately, these worthwhile attributes are often outweighed by the 10 disadvantage that online information delivery makes it relatively easy to access pristine digital content and to pirate the content at the expense and harm of the publisher.

Piracy of online digital content is not yet a great problem. Most premium content that is available on the Web is of low value and therefore casual and organized pirates do not yet see an attractive business stealing and reselling content. Increasingly, higher-value content is becoming available. Audio 15 recordings are available now, and as bandwidths increase, video content will start to appear. With the increase in value of online digital content, the attractiveness of organized and casual theft increases.

The unusual property of digital content is that the publisher or reseller transmits the content to a client, but continues to restrict rights to use the content 20 even after the content is under the sole physical control of the client. For instance, a publisher will often retain copyright to a work so that the client cannot reproduce or publish the work without permission. A publisher could also adjust pricing according to whether the client is allowed to make a persistent copy, or is just allowed to view the content online as it is delivered. These scenarios reveal a 25 peculiar arrangement. The user that possesses the digital bits often does not have full rights to their use; instead, the provider retains at least some of the rights. In a very real sense, the legitimate user of a computer can be an adversary of the data or content provider.

“Digital rights management” is fast becoming a central theme as online commerce continues its rapid growth. Content providers and the computer industry must quickly address technologies and protocols for ensuring that digital data is properly handled in accordance with the rights granted by the publisher. If 5 measures are not taken, traditional content providers may be put out of business by widespread theft or, more likely, will refuse to deliver content online.

Traditional security systems ill serve this problem. There are highly secure schemes for encrypting data on networks, authenticating users, revoking users, and storing data securely. Unfortunately, none of these systems address the assurance 10 of content security after it has been delivered to a client’s machine. Traditional uses of smart cards offer little help. Smart cards merely provide authentication, storage, and encryption capabilities. Ultimately, useful content must be delivered to the host machine for display, and again, at this point the bits are subject to theft. Cryptographic coprocessors provide higher-performance smart-card services, and 15 are usually programmable; but again, any operating system or process, trusted or not, can use the services of the cryptographic processor.

There appear to be three solutions to this problem. One solution is to do away with general-purpose computing devices and use special-purpose tamper-resistant boxes for delivery, storage, and display of secure content. This is the approach 20 adopted by the cable industry and their set-top boxes, and appears to be the model for DVD-video presentation. The second solution is to use proprietary data formats and applications software, or to use tamper-resistant software containers. The third solution is to modify the general-purpose computer to support a general model of client-side content security and digital rights management.

25 This invention is directed to a system and methodology that employs the third category of solutions.

The fundamental building block for client-side content security is a secure operating system. If a computer can be booted into an operating system that is

trusted to honor content rights, and only allows authorized applications to access rights-restricted data, then data integrity within the machine can be assured. The stepping-stone to a secure operating system is sometimes called “Secure Boot” If secure boot cannot be assured, whatever rights management system the OS 5 provides can always be subverted by booting into an insecure operating system.

Secure boot of an operating system is usually a multi-stage process. A securely booted computer runs a trusted program at startup. The trusted program loads another program and checks its integrity, e.g., by using a code signature, before allowing it to run. This program in turn loads and checks subsequent 10 layers. This proceeds all the way to loading trusted device drivers, and finally a trusted application. Related patent application MS1-282US describes an overall method of securely booting an operating system, and also notes related technology.

Booting an operating system or other program securely requires some way to 15 execute code such that the code cannot be tampered with as it is being executed, even by one who is in physical possession of the computer that executes the code. In the scenarios discussed above, digital content is loaded from a network or from a medium into a personal computer at a remote location. The PCs’ owners have 20 full freedom to run arbitrary programs for compromising any safeguards, to replace ROM containing trusted BIOS code, to bypass dongles, to introduce rogue hardware, even to analyze signals on buses. Today’s low-end computers are open systems, both logically and physically. Indeed, most computers of all kinds are open, at least to those having supervisory privileges and physical possession.

At the same time, conventional techniques for restricting subversion in this 25 environment impose either unacceptable burdens upon legitimate users or they are unacceptably expensive. S. T. Kent’s Ph.D. thesis, “Protecting Externally Supplied Software in Small Computers”, MIT Laboratory for Computer Science 1980, is an early proposal for tamper-resistant modules. S. R. White, “ABYSS: A

Trusted Architecture for Software Protection”, Proceedings, 1987 IEEE Symposium on Security and Privacy, pp. 38-51, presents a trusted architecture having a secure processor in a tamper-resistant package such as a chip, for enforcing limitations to execute application code. This system, however, would 5 require major changes to existing processor architectures, and would still be limited to the small instruction set of a primitive security coprocessor. Also, it is limited to on-board, physically inaccessible memory dedicated to security functions.

The practicality of trusted operating systems still requires an inexpensive way 10 to execute code that cannot be easily modified or subverted, a way that does not necessitate new or highly customized processors and a way that performs as much as possible of the secure execution in software.

Summary of the Invention

15 The present invention provides a more general-purpose microprocessor and memory-system architecture that can support authenticated operation, including authenticated booting of an operating system. This new class of secure operation is called curtained execution, because it can be curtained off and hidden from the normal operation of the system. The code executed during such operation is called 20 curtained code; it can preserve secret information even from a legitimate user in physical possession of an open computer.

The invention allows users to load and reload data and programs for authenticating operations without physically modifying (or having someone else modify) their computers. For example, a software or content provider can provide 25 encrypted keys along with code for manipulating those keys to users without fear of compromising the keys, because the code can only be executed in a manner that preserves their secrecy.

Curtained operation does not make great demands upon a processor, and requires few modifications from standard designs. It allows innovation in particular implementations and applications to take place at software-development cycle times, rather than at the slower pace of hardware versions. It gives content providers and program developers an opportunity to design and personalize secure operations for their specific needs. Further, curtained code is not limited to the small instruction sets, program sizes or memory requirements of dedicated secure processors or coprocessors, and it promises applications beyond its core purpose of authenticating other programs.

Curtained operation generalizes the concept that certain memory regions are only accessible to certain code. Whereas conventional memory-protection schemes grant or deny memory-access rights to designated address ranges based upon an internal kernel or supervisory state of the processor regardless of the code executing, curtained operation ties access rights to certain code. Curtained code can only be executed from certain locations, and the physical address from which it is executed determines its access rights. Other applications or operating system code does not have the necessary rights to modify the curtained memory regions or to obtain secrets stored in such regions.

Curtained execution also forces atomic execution of the curtained code, to prevent spurious code from hijacking its operation or from stealing secret information stored in machine registers following a legitimate initial call.

The Drawing

Fig. 1 is a block diagram of a computer system in which the invention can be implemented.

Fig. 2 is a symbolic map showing memory regions organized according to the invention.

Fig. 3 is a block diagram of a processor for carrying out the invention.

Fig. 4 is a flowchart of a method for curtained code execution according to the invention.

5

Detailed Description

This description and the accompanying drawing illustrates specific examples of embodiments in which the present invention can be practiced, in enough detail to allow those skilled in the art to understand and practice the invention. Other embodiments, including logical, electrical, and mechanical variations, are within the skill of the art, as are other advantages and features of the invention not explicitly described. The scope of the invention is to be defined only by the appended claims, and not by the specific embodiments described below.

The description proceeds from an illustrative environment to an organization for a secure memory area and then to mechanisms for executing trusted code that can access the memory. Finally, some representative applications of curtained operation are presented.

Environment

Fig. 1 is a high-level diagram of an illustrative environment 100 having software 110 and hardware 120 for hosting the invention as executable instructions, data, and/or electronic and mechanical components. Other suitable environments, and variations of the described environment are also possible.

Hardware components 120 are shown as a conventional personal computer (PC) including a number of components coupled together by one or more system buses 121 for carrying instructions, data, and control signals. These buses may assume a number of forms, such as the conventional ISA, PCI, and AGP buses.

Some or all of the units coupled to a bus can act as a bus master for initiating transfers to other units. Processing unit 130 may have one or more microprocessors 131 driven by system clock 132 and coupled to one or more buses 121 by controllers 133. Internal memory system 140 supplies instructions and data to processing unit 130. High-speed RAM 141 stores any or all of the elements of software 110. ROM 142 commonly stores basic input/output system (BIOS) software for starting PC 120 and for controlling low-level operations among its components. Bulk storage subsystem 150 stores one or more elements of software 110. Hard disk drive 151 stores software 110 in a nonvolatile form.

5 Drives 152 read and write software on removable media such as magnetic diskette 153 and optical disc 154. Other technologies for bulk storage are also known in the art. Adapters 155 couple the storage devices to system buses 121, and sometimes to each other directly. Other hardware units and adapters, indicated generally at 160, may perform specialized functions such as data encryption,

10 signal processing, and the like, under the control of the processor or another unit on the buses.

15

Input/output (I/O) subsystem 170 has a number of specialized adapters 171 for connecting PC 120 to external devices for interfacing with a user. A monitor 172 creates a visual display of graphic data in any of several known forms. Speakers 173 output audio data that may arrive at an adapter 171 as digital wave samples, musical-instrument digital interface (MIDI) streams, or other formats. Keyboard 174 accepts keystrokes from the user. A mouse or other pointing device 175 indicates where a user action is to occur. Block 176 represents other input and/or output devices, such as a small camera or microphone for converting video and

20 audio input signals into digital data. Other input and output devices, such as printers and scanners commonly connect to standardized ports 177. These ports include parallel, serial, SCSI, USB, FireWire, and other conventional forms.

25

Personal computers frequently connect to other computers in networks. For example, local area network (LAN) 180 connect PC 120 to other PCs 120' and/or to remote servers 181 through a network adapter 182 in PC 120, using a standard protocol such as Ethernet or token-ring. Although Fig. 1 shows a physical cable 183 for interconnecting the LAN, wireless, optical, and other technologies are also available. Other networks, such as wide-area network (WAN) 190 can also interconnect PCs 120 and 120', and even servers 181, to remote computers 191. Fig. 1 illustrates a communications facility 192 such as a public switched telephone network for a WAN 190 such as an intranet or the internet. PC 120 can employ an internal or external modem 193 coupled to serial port 177. Other technologies such as packet-switching ISDN, ATM, DSL, frame-relay are also available. In a networked or distributed-computing environment, some of the software 110 may be stored on the other peer PCs 120', or on computers 181 and 191, each of which has its own storage devices and media.

Software elements 110 may be divided into a number of types whose designations overlap to some degree. For example, the previously mentioned BIOS sometimes includes high-level routines or programs which might also be classified as part of an operating system (OS) in other settings. The major purpose of OS 111 is to provide a software environment for executing application programs 112 and for managing the resources of system 100. An OS such as Microsoft® Windows® or Windows NT® commonly implements high-level application-program interfaces (APIs), file systems, communications protocols, input/output data conversions, and other functions.

Application programs 112 perform more direct functions for the user. A user normally calls them explicitly, although they can execute implicitly in connection with other applications or by association with particular data files or types. Modules 113 are packages of executable instructions and data which may perform functions for OSs 111 or for applications 112. Dynamic link libraries (.dll) and

class definitions, for instance, supply functions to one or more programs. Content 114 includes digital data such as movies, music, and other media presentations that third parties make available on media or by download for use in computer 120. This material is frequently licensed for a charge, and has certain restrictions placed 5 upon its use.

Secure Memory Organization

Fig. 2 is a symbolic map of a memory space 200 in system 100. For purposes of illustration, consider it to have a potential size of 4Gbytes, so that 32 bits of 10 address suffice to access all of it. Space 200 can exist in a single physical memory, or in several different kinds of storage, such as ROM, read/write RAM, flash RAM, and so forth. Also, partially or totally separate address spaces are a straightforward extension. Space 200 has three hierarchical rings 210, 220, and 230 relevant to the present discussion. Although the information stored in these 15 rings can be similar to that contained in the rings sometimes used in processors that employ conventional privilege levels or operational modes, their mechanism differs.

Ring 210 is called Ring C or the outer ring, and has no protection or security against any kind of read or write access by any code located there or in the other 20 rings in the present system, and normally occupies almost all of the available address space. All normal user code and data resides in this ring. The operating system, including the kernel, also resides there. Ring C has no read or write access to the other two rings.

The secure rings 220 and 230 together comprise the secure or curtained region 25 of memory. No program code in Ring C has any access to data within them. Ring C code, can, however, be provided some ability to initiate the execution of code located there, as described below. Conversely, any code in rings 220 and 230 has

full access to Ring C, including reading and writing data, and executing program code.

Secure ring 220, also called Ring B, is an inner ring to Ring C, and has full access privileges to its outer Ring C; but Ring B is in turn an outer ring with respect to ring A, and thus has only restricted access to this inner ring. In this embodiment, the major purpose of Ring B is to hold most of the code that carries out authenticated-boot operations as mentioned above and in Application docket MS1-282US. Thus, it can have both semipermanent storage such as nonvolatile flash RAM for code routines and volatile read/write memory for temporary data such as keys. A megabyte or less of the total address range would likely suffice for Ring B.

Secure ring 230, also called Ring A is an inner ring to both Rings B and C, and has full access to them for both code and data. It can also employ both nonvolatile and volatile technologies for storing code and data respectively. Its purpose in this embodiment is to store short loader and verifier programs and keys for authentication and encryption. Under the proper conditions, this code and data can be loaded in the clear. The address space required by Ring A is generally much smaller than that of Ring B. That is, this exemplary embodiment has the Ring A address range within the address range of Ring B, which in turn lies within the address range of Ring C. The address ranges of the rings need not be contiguous or lie in a single block. In order to prevent the access restrictions of the curtained rings from being mapped away by a processor, the address ranges of Rings A and B can be treated as physical addresses only. In one embodiment, virtual addresses are conventionally translated into their corresponding real addresses, and then the restrictions are interposed at the level of the resulting real addresses. Alternatively, a mechanisms could disable virtual addressing when certain addresses are accessed.

In the contemplated area of authentication of rights, it can be desirable to allow multiple parties to emplace their own separate authentication code and data that cannot be accessed by any of the other parties. For example, the manufacturer of the processor, the provider of the operating system or trusted application programs, and certain organizations that furnish digital content may all desire to execute their own authentication or other security routines and manage their own keys. At the same time, each party should be able to use code and data in the unsecure outermost Ring C, and to execute certain routines in the innermost Ring A. Dividing Ring B into peer subrings 221, 222, and 223 permits this type of operation. Ring 221, called Subring B1, has the privileges and restrictions of Ring B, except that it cannot access subring 222 or 223. It can access any part of Ring B that lies outside the other subrings, however. In this way, Subring B1 can function as though it were the only middle ring between Rings A and C for some purposes. Rings 222 (Subring B2), and 223 (Subring B3) operate in the same manner. A typical PC-based system might have three or four subrings, of 64-128KBytes each. The code in these subrings is normally updated seldom, so that conventional flash memory is a convenient technology. Alternatively, the Ring-A loader could load the code and keys into RAM from an encrypted storage on disk on demand. Each subring will also require a small amount of scratch RAM, although rewritable flash memory might be suitable here as well; it might be desirable to use this for persisting the state of the system after a reboot. For extra flexibility, the memory available to the curtained memory subsystem can be allocated under the control of the Ring-A executive code. In order that no untrusted party can manipulate the memory map to reveal secrets, the map of the subrings in the Ring-B memory is kept in flash storage in curtained memory, under control of the curtained-memory controller in ring A.

In presently contemplated authentication procedures, Ring A code and keys are loaded under conditions in which protection against snoopers is not necessary;

for example, they can be loaded when the microprocessor is manufactured. This simple step eliminates any requirement for building any cryptographic capabilities into the processor itself. Accordingly, Ring A code and keys can be stored in permanent ROM, with only a few hundred bytes of scratchpad RAM. This Ring A 5 code is designed to load further curtained code and keys into ring B memory segments through a physically insecure channel, such as a public network, in such a manner that an eavesdropper, including even the owner of the target computer, cannot discover any secret information contained therein. This downloaded code, 10 operating from the secure memory, then performs the authentication operations that third parties require before they will trust their valuable content to the rights-management software of the system. This new bootstrapping procedure permits building a wide class of secure operations and associated secret keys with greater security than would be possible in traditional assembly code, even with some form of authentication routines.

15 However, there are no restrictions on the code that can be loaded into any of the Ring-B memory areas. Examples of Ring-B code include smartcard-like applications for key management, secure storage, signing, and authentication. Further examples include electronic cash storage, a secure interpreter for executing encrypted code, modules for providing a software licenses necessary for 20 a piece of software to run. It is also possible to load only a part of an application, such as a module that communicates with a media player in unsecure memory for reducing software piracy.

Executing Curtained Code

25 The foregoing shows how untrusted code can be prevented from accessing the contents of a secure memory. The trusted code that is permitted to perform secure operations and to handle secret data is called curtained code. In other systems, such code must be executed within a privileged operating mode of the processor not accessible to the user, or from a separate secure processor. In the present

invention, however, curtained code can only be executed from particular locations in memory. If this memory is made secure against intrusion, then the curtained code can be trusted by third parties. Other features restrict subversion through attempts at partial or modified execution of the curtained code.

5 Fig. 3 is a block diagram showing relevant parts of a microprocessor 300 that can serve as part or all of processing unit 131, Fig. 1. Internal buses 301 carry data, address, and control signals to the other components of the processor on the integrated-circuit chip or module. Line 302 carries some of these signals to and from bus controller 133. Conventional function or execution units 310 perform 10 operations on data from external memory, from register files 320, from cache 330, from internal addressable memory 340, or from any other conventional source. Memory 340, located on the same chip or module as the rest of processor 300, can have a number of technologies or combinations of technologies, such as dynamic 15 read/write, read-only, and nonvolatile such as flash. The internal memory in this implementation partakes of the same address sequence as external system memory 140, although it can have or be a part of another sequence. The curtained memory rings can be partly or totally contained in addresses located within memory 340.

Control unit 350 carries out a number of operations for sequencing the flow of 20 instructions and data throughout the processor; line 304 symbolizes control signals sent to all of the other components. Interrupt logic 351 receives interrupt requests and sends system responses via lines 305; in some systems, interrupt logic is conceptually and/or physically a part of controller 133. A conventional instruction 25 pointer holds the address of the currently executing instruction. Instruction decoder 353 receives the instruction at this address on line 306, and produces a sequence of control signals 304 for executing various phases of the instruction. In modern pipelined and superscalar microprocessors, blocks 352 and 353 become very complex as many instructions are in process at the same time. Their basic functions, however, remain the same for the present purpose.

Control unit 350 further includes a specification or map 354 of one or more address ranges of the memory addresses desired to be curtained. The specification can be in any desired form, such as logic circuitry, a read-only table of addresses or extents, or even a small writable or rewritable storage array. If the addresses 5 are in memories having separate address sequences, additional data specifying the particular memories can be added to the addresses within each sequence. A detector or comparator 355 receives the contents of instruction pointer 352 and the curtained-memory map 354. A curtained memory having multiple rings, subrings, or other levels can have a separate specification for each of the curtained regions. 10 Alternatively, a single specification can explicitly designate the ring or subring that each address range in the specification belongs to.

If the current instruction address from pointer 352 matches any of the addresses in map 354, that instruction is included in a particular curtained code ring or module. Curtain logic 356 then permits the control unit to issue signals 15 304 for performing certain operations, including reading and writing memory locations in the same ring, or a less privileged ring that might contain secrets. Additionally, as described below, certain opcodes are restricted to executing only when the CPU is executing curtained code. For example, if decoder 353 is executing an instruction not located within the range of curtained memory, and if 20 that instruction includes an operand address located within the curtained-memory specification, control unit 350 blocks the signals 304 for reading the data at that address and for writing anything to that address. If a non-privileged access is attempted, the CPU or memory system can flag an error, fail silently, or take other appropriate action. If it is desired to place the curtain logic on a chip other than 25 the processor, a new microprocessor instruction or operating mode can strobe the instruction pointer's contents onto an external bus for comparison with the curtained address ranges.

The execution of trusted code routines is frequently initiated by other programs that are less trusted. Therefore, curtain logic 356 must provide for some form of execution access to the curtained code stored in Rings A and B. However, full call or jump accesses from arbitrary outside code, or into arbitrary locations of 5 the curtained memory regions, might possibly manipulate the secure code, or pieces of it, in a way that would reveal secret data or algorithms in the curtained memory. For this reason, logic 356 restricts execution entry points into curtained memory regions 220 and 230 as well as restricting read/write access to those regions. In one embodiment, the curtained code exposes certain entry points that 10 the code writers have identified as being safe. These often occur along functional lines. For instance, each operation that a piece of curtained code can perform has an accompanying entry point. Calling subroutines at these entry points is permitted, but attempts to jump or call code at other entry points causes an execution fault.

15 An alternative allows automated checking of entry points and provides additional granularity of rights by permitting entry to curtained memory functions only through a special entry instruction. For example, a new curtained-call instruction, CCALL Ring, Subring, OpIndex, has operands that specify a ring, a subring, and a designation of an operation whose code is located within that 20 ring and subring. This instruction performs conventional subroutine-call operations such as pushing a return address on a stack and saving state information. The stack or the caller's memory can be used to pass any required parameters. A conventional RETURN instruction within the curtained code returns control to the calling routine. Return values can be placed in memory, registers, 25 etc.

When decoder 353 receives a CCALL instruction, curtain entry logic 356 determines whether the calling user code has the proper privileges, and whether the instruction's parameters are valid. If both of these conditions obtain, then the

instruction is executed and the curtained routine is executed from its memory ring. If either condition does not hold, logic 356 fails the operation without executing the called code.

Logic 356 determines whether or not to execute the code by comparing the privilege level of the calling code and the operation-index parameter, and potentially whether the processor is already executing some other curtained code, with entries in a jump-target table 357 stored in a location accessible to it. The logic to enforce these requirements can be implemented in the memory controller 356, or by code executing in a highly privileged ring such as Ring A. Table I below illustrates one form of jump-target table. The table can be stored in the same curtained memory block as the code itself, or in a memory block that is more privileged; or it can be stored in special-purpose storage internal to the CPU or memory manager.

15

Table I

Index	Target Address	User	Kernel	Curtain
0	BAB-PC	FALSE	TRUE	TRUE
1	REVEAL-PC	TRUE	TRUE	TRUE
2	LOAD-PC	FALSE	FALSE	TRUE

An entry for each index, 0-2, gives the (symbolic) target or start address of the code for that operation, and the privileges levels—user, kernel, or curtained—that are permitted to execute the code. “Curtained” level means that only other curtained code can call the routine. Other or finer privilege levels are possible. As an alternative to the above jump table, entry logic 356 could permit only a

single entry point into each ring of curtained memory, and employ a passed parameter to specify a particular operation. Or it could, for example, permit calls only to addresses that are predefined as the beginnings of operations. The curtained code itself could verify and call the operation.

- 5 Restricting call access to curtained code within processor 300 still leaves open the possibility that outside rogue programs or devices might be able to hijack the code after its execution has begun in order to obtain secrets left in registers, or to otherwise modify machine state to subvert operation. Therefore, control unit 350 must ensure atomicity in executing the curtained code: once started, the code must
- 10 perform its entire operation without interruption from any point outside the secure curtained-memory regions. In many cases, it is not necessary to execute an entire function atomically, but only a part. For example, only the code that verifies a bus-master card's identity need be performed atomically, and not its total initialization module.
- 15 Modern, open computer systems present a number of paths for obtaining access to any hardware, software, and data within the system. Personal computers in particular have been designed with very little thought for security, and with even less provision for restrictions against their legitimate users. Because many advantages of PCs and similar systems flow from an open environment, however,
- 20 the protection for atomicity should impose as few restrictions as possible. The following outlines the major forms of gaining access to a memory in a conventional PC, and some of the ways to prevent access to a curtained segment of memory. Different systems may employ different combinations of these and other access restrictions.
- 25 Interrupts offer almost unlimited access to system resources. A simple way to prevent an interrupt from subverting curtained code is to issue a privileged instruction that causes a microprocessor to switch off all interrupts until a companion instruction switches them back on. A new instruction such as

SnoopIntrerrupts Ring, Subring, OpIndex can call a curtained operation instead of the requested interrupt routine when an interrupt tries to access memory in a designated ring or subring, or operation. ¹This can also be managed by having the curtained code set up the interrupt handlers to execute

5 trusted curtained code. However, it is still important to protect the entry point into the curtained operation (that sets the interrupt vector) itself be protected against interruption so that the interrupt mechanism can be subverted by a malicious program.

An instruction having the form SetOpaque MemoryStart ,

10 MemoryLength / SetInterruptThrowError / SetTransparent does not switch off interrupts, but rather modifies the microprocessor" behavior. When an interrupt occurs, the processor clears all registers, except the stack pointer, before the interrupt is fielded. It is useful for long-running curtained operations that could reveal sensitive information, such as partial keys, if they

15 were interrupted. An operand of this instruction can specify a memory range that the processor also clears before the interrupt is serviced. The first switch of the instruction activates a variant that causes a processor fault when an interrupt occurs, even in user mode. The user code can then disable operations and process—or decide not to process—the interrupt. The second switch turns off the

20 SetOpaque execution mode. These can be user-mode operations, if desired. In at least some circumstances, this instruction should fault the processor when returning from the interrupt, to prevent an undesired jump into the middle of curtained code that might have been executing when the interrupt took control.

Illegal-operation and page faults are commonly encountered types of

25 interrupt. Some systems might wish to handle these interrupts in the normal manner, and to disable only those interrupts generated asynchronously or externally to the microprocessor. Faults or interrupts produced by debuggers,

however should be disabled; one of the oldest and easiest ways to hack any code is to pry it open with a debugger.

System buses commonly allow devices other than the processor to access memory on them. Bus master cards in a PC, for example, have the ability to read and write main memory. Curtained memory in this environment may require restrictions upon bus access to memory modules. If the secure memory is located on the same chip as the microprocessor, or within the same physically secure module, merely causing the processor not to relinquish the bus during curtained operation may offer adequate protection. Most cases of interest here, however, must assume a trusted chipset, and will protect the bus via a controller such as 133, Fig. 1. Block 303 in Fig. 3 represents one possible location for the memory-bus lock.

A new privileged instruction, `LockBus`, can disable all accesses to memory apart from those initiated by the processor executing authorized code. A 15 companion `UnlockBus` instruction terminates this mode. In most systems, these instructions should be executable only in a privileged mode. An alternative type of instruction detects memory reads and writes by all devices on the bus other than the processor. A simple `SnoopBus [Throw]` form can set a flag, cause a fault, clear certain registers and/or memory, or call a curtained operation to cancel any 20 outstanding privileges or identity. Parameters such as `Ring`, `Subring`, `OpIndex` can specify one or more memory ranges, thus allowing multiple processors and bus-master controllers to continue operating. Parameters such as `MemoryStart`, `MemoryLength` can monitor bus requests from other bus agents, then zero out a memory block before relinquishing the bus to the other 25 agents. Any method of destroying the contents of a memory or register can obviously be used instead of zeroing. This type of instruction could be useful for user-mode application programs to protect their curtained operations from prying by the operating system or by debuggers, and might be allowed in user code.

Another limitation available in some environments is to restrict outside devices only until a trusted routine has verified them or initialized them properly.

One further hardware restriction that is valuable from the perspective of protection against a computer's expansion cards is the ability to disable all DMA or bus-mastering activity from a device plugged into a particular PC slot until the device is explicitly identified, initialized and made safe. Early in the boot sequence, all bus-master activity is disabled on the PC bus controller: the slots are locked. The devices are identified and initialized using a conventional type of programmed IO. Only after correct initialization are the slots unlocked one by one, so that full functionality is available. Devices that are unknown, or that do not behave as they should, will not be enabled, and hence can not subvert operation or steal secrets. This action is called "slot locking."

Fig. 4 is a flowchart 400 of a method for providing curtained execution protection in a processor such as 300. For a single-level curtained memory, method 400 refers to the entire curtained region. For a memory organization such as 200 having multiple rings or subrings, the term "curtained region" means the levels inside or beside the ring in which the current instruction is located. For example, the curtained region for an instruction whose address is in Ring C in Fig. 2 comprises Ring B (including all its subrings) and Ring A; the curtained region for an instruction in Subring B1 comprises Subrings B2 and B3 (but not the rest of Ring B) and Ring A.

After block 410 decodes the current instruction, blocks 420 test memory addresses associated with the instruction. If the instruction uses virtual addresses, tests 420 operate upon the physical addresses as translated by decoder 410. Block 25 421 determines whether the instruction accesses any memory location during its execution. An instruction might read an operand or write data to a memory address, for example. If the instruction does not access any memory, or at least any memory that might contain a curtained region, then block 430 executes the

instruction. If the instruction does involve a memory location, block 422 tests the address to determine whether it is within a region that is curtained off from the current region. If not, block 430 executes the instruction. If so, block 423 asks what type of access the instruction requests. If the access is anything other than 5 the special curtained-call opcode, then block 440 signals a fault, and an appropriate error routine or logic circuit blocks the access. Other accesses include reading data from the location, writing data to it, or executing a normal instruction there.

The only access permitted into a curtained-memory ring is an execution access 10 by a particular kind of instruction, such as the curtained call (CCALL) discussed above. If block 423 detects that this instruction desires to initiate execution of code at a location inside a region curtained from the current region, block 424 determines whether the target entry point is valid—that is, whether the requested index is in the jump table. Block 425 then determines whether the current 15 instruction has the privilege level required to invoke the operation at the desired location. If either test fails, block 440 produces a fault. If both pass, block 450 executes the curtain-call instruction as described above.

Blocks 460 navigate among the rings and subrings of the curtained memory. A CCALL instruction causes block 461 to open the curtained-memory ring 20 containing the target address of the call. That is, it makes that ring the current ring for the purposes of method 400. A routine starting at that address thus has read/write and execution access to the memory of the ring, and only rings inside or peer to that ring are now restricted curtained memory. Block 461 also engages any extra protection for ensuring atomicity of the routine being executed at the 25 new current level, such as interrupt suspension or bus locking. A routine executing in curtained memory can end with a normal Return instruction. If the routine was called from a less secure ring, block 462 causes block 463 to close the

current ring and retreat to the ring from which the call was made, either a less secure ring of curtained memory, or the outer, unsecured memory of Ring C.

Operations with Curtained Code

5 The following illustrate a few representative applications of curtained operation.

10 Loading and reloading secure routines is difficult in conventional practice. The procedure below allows even an untrusted user to field-load curtained code and secret keys into Ring-B memory without being able to discover the secret 10 keys.

15 (1) Execute an authentication and key-exchange protocol. If the protocol runs successfully, keep the session key in Ring-A curtained RAM. The authentication step assures the software and key publisher that the target is truly a piece of trusted code executing in a protected environment, and not an arbitrary application or operating system that will read and abuse the secret keys. This can be assured by equipping the loader with a public-key cryptography key pair and a certificate from a trusted manufacturer or publisher that indicates the source of the code, and therefore that it is executing in a trusted 20 (curtained) computer system.

25 (2) After successful completion of the previous step, load a block of encrypted code and accompanying keys into Ring-B flash memory. This code should be protected from alteration by a checksum, digital signature, or other means. It should be preceded by any entry-protection mechanism, such as a jump table.

(3) Verify Ring B, if secret keys are to be granted to this ring, or if desired for extra assurance that the code has not been tampered with.

5

Verification can be carried out by generating a signature with a secret key such as a hash digest of all the code in one of the B rings and a nonce supplied with the code to be loaded in this session. A nonce is a single-use unpredictable value, of the type used in a zero-knowledge proof.

10

(4) Get the name or description of a set of opcodes in a Ring-B code-set. This allows a user to select a set of curtained operations—say from the processor-manufacturer’s subring—that the preceding step has verified. This name permits an application or operating system to select an appropriate code-set among the many that might be loaded in a system. It could be a simple textual description such as “MS Key Store 3.0,” or it could be a cryptographic digest of the data that comprises the curtained memory region.

15

(5) Swap codes sets to and from secure memory or to and from some other persistent store such as disk, if there are more code sets than Ring-B slots available for them. If a ring-B module contains secret keys, the Ring-A loader must encrypt them prior to exporting them to main memory. In most cases the operating system handles the transfer to and from disk.

20

(6) Allocate specific Ring-B memory to particular code sets, and collect garbage to avoid holes.

The following boot-block pseudocode sets an identity to the public key of a piece of signed code.

25

```
RetryLabel:  
    CCALL BeginBoot  
    [MAC, Signature, Public Key]    // of all of  
    bootblock
```

```
[check signature of next code and data block] [PKs  
of next blocks]  
    if(SignatureOK) CCAL CompleteBoot  
    else CCAL TerminateBoot  
5      [next section of boot code]
```

The three curtained-code operations for setting this identity are:

```
[User=FALSE, Kernel=True, Curtained=TRUE]  
BeginBoot  
OldStackPointer=StackPointer  
10     SetOpaque  
TemporaryIdentity=NULL  
SnoopBus TerminateBoot  
SnoopInterrupts TerminateBoot  
Calculate MAC of bootblock from address inferred  
15     from *SP
```

```
If(signature good for stated public key)  
TempId=PublicKey  
Else TempId=PublicKey  
[Zero registers and scratch RAM]  
20     Return
```

```
[User=TRUE, Kernel=True, Curtained=TRUE]  
TerminateBoot  
TempId=NULL  
25     StackPointer=OldStackPointer+1  
Goto ReTryLabel
```

```
[User=FALSE, Kernel=True, Curtained=TRUE]  
CompleteBoot
```

```
CodeIdentity=TempId
UnSnoopInterrupts
UnSnoopBus
SetTransparent
```

5 Given a seed and a processor identity, the next code swatch generates a storage key for securing content. The seed and the return value are stored in the calling program's memory space.

```
[User=FALSE, Kernel=True, Curtained=TRUE]
GenerateKey (&InSeed, &ReturnVal)
10 SetOpaque
If (CodeIdentity==NULL) return NULL
[Compute a pseudo random number 'Key' using a seed
derived from InSeed, MySecretKey, codeIdentity]
RetVal=Key
15 [zero registers and scratch RAM]
SetTransparent
Return
```

Checking OS identity is a major application for curtained operation. The first time the following operation executes, it builds a digest of the OS. Later 20 invocations check new digests against the first one to ensure that the OS image has not changed, and revokes its identity if it has.

```
[User=FALSE, Kernel=True, Curtained=TRUE]
CheckIdentity (MemoryTable)
SetOpaque
25 NewDigest=[CreateDigest]
If (OldDigest==NewDigest) {
    SetTransparent
    Return
}
```

```
1 If (OldDigest!=NewDigest) {  
2     codeIdentity=NULL  
3     SetTransparent  
4     Return  
5 }  
6 SetTransparent  
7 Return
```

The initial identity can be derived from other steps, or built up in stages in
10 curtained RAM before newly loaded code is executed. Transitive trust then
ensures that security is as good as the initial check.

Conclusion

The foregoing describes a system and method for curtained execution of code
that can be trusted by a third party in an environment where a possibly hostile
15 person has physical possession of the system upon which the trusted code
executes. It permits field loading of sensitive code and data by such a person.
Other advantages and variations will be apparent to those skilled in the art.

For example, different security requirements and different systems may permit
different or relaxed provisions for securing the curtained memory and code against
20 certain kinds or levels of attack. For example, legacy systems might not permit all
of the components described above to be fabricated in a single chip. In this case, a
potted or otherwise secure chipset between the existing microprocessor and its
motherboard socket can implement curtained execution and memory. Some
existing microprocessors have system management or other restricted operating
25 modes that can provide some or most of the security requirements. Curtained
operation can be extended to additional rings; all or most of the operating system
might be placed in a curtained ring, for example.

Dynamic resizing or layout of secure memory ring is feasible in some cases; the curtain logic or memory manager should clear out ring contents and memory pages before their access rights are changed. Although the present implementation permits only real addresses in curtained memory, virtual addressing may be

5 feasible, given adequate safeguards against mapping away the access security.

Some processors already possess system management modes that provide access, entry-point, and atomicity restrictions that may provide enough security that curtained memory could be mapped into their address spaces, especially if only a single curtained ring or region is needed.

10 Other applications for curtained operation can be easily imagined. A secure interpreter for encrypted code can be executed from curtained memory. Certified execution can construct a hashed digest of actual executed code that is attested as correct by curtained code. In addition to authenticating an OS upon boot-up, calls for private keys can be made to require a curtained operation to check its

15 continuing integrity. Where rights are given for a fixed number of iterations or for a certain time interval, curtained code can implement a monotonic counter or clock. Certificate revocation lists, naming components that are known to be compromised or otherwise undesirable, can employ such a secure counter to prevent components from being removed from a list. A number of rights-
20 management functions demand a tamper-resistant log. A signed or encrypted Ring-C file having a Ring-B digest or key can serve this purpose. Secure interpretation of a certificate that grants rights to code identity enables more levels of indirection between boot-code authentication and rights to content; this facilitates fixing bugs and updating components without losing keys already stored
25 in a system. Any rights that rely upon continued secrecy of keys or the strength of particular cryptographic algorithms is fragile. Curtained operation is sufficiently flexible to field-load changes to circumvent compromises of secret data or code.

A Ring-B subring can also provide smart-card types of service, and could offer those services to a trusted operating system.